

Visual Pivoting for (Unsupervised) Entity Alignment

Fangyu Liu¹, Muhao Chen^{2,3}, Dan Roth², Nigel Collier¹

¹ Language Technology Lab, TAL, University of Cambridge, UK

² Department of Computer and Information Science, University of Pennsylvania, USA

³ Viterbi School of Engineering, University of Southern California, USA
{fl399, nhc30}@cam.ac.uk, muhaoche@usc.edu, danroth@seas.upenn.edu

Abstract

This work studies the use of visual semantic representations to align entities in heterogeneous knowledge graphs (KGs). Images are natural components of many existing KGs. By combining visual knowledge with other auxiliary information, we show that the proposed new approach, **EVA**, creates a holistic entity representation that provides strong signals for cross-graph entity alignment. Besides, previous entity alignment methods require human labelled seed alignment, restricting availability. **EVA** provides a completely unsupervised solution by leveraging the visual similarity of entities to create an initial seed dictionary (visual pivots). Experiments on benchmark data sets DBP15k and DWY15k show that **EVA** offers state-of-the-art performance on both monolingual and cross-lingual entity alignment tasks. Furthermore, we discover that images are particularly useful to align long-tail KG entities, which inherently lack the structural contexts that are necessary for capturing the correspondences. Code release: <https://github.com/cambridgeltl/eva>; project page: http://cogcomp.org/page/publication_view/927.

1 Introduction

Knowledge graphs (KGs) such as DBpedia (Lehmann et al. 2015), YAGO (Rebele et al. 2016) and Freebase (Bollacker et al. 2008) store structured knowledge that is crucial to numerous knowledge-driven applications including question answering (Cui et al. 2017), entity linking (Radhakrishnan, Talukdar, and Varma 2018), text generation (Koncel-Kedziorski et al. 2019) and information extraction (Hoffmann et al. 2011). However, most KGs are independently extracted from separate sources, or contributed by speakers of one language, therefore limiting the coverage of knowledge. It is important to match and synchronise the independently built KGs and seek to provide NLP systems the benefit of complementary information contained in different KGs (Bleiholder and Naumann 2009; Bryl and Bizer 2014). To remedy this problem, the Entity Alignment (EA)¹ task aims at building cross-graph mappings to match entities having the same real-world identities, therefore integrating knowledge from different sources into a common space.

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹The *entity* in EA refers to real-world objects and concepts.

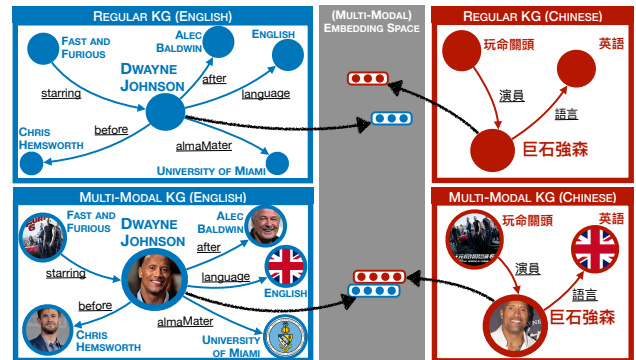


Figure 1: A challenging example which heavily benefits from using vision to align entities between cross-lingual KGs in DBP15k. We display the neighbourhoods of entity DWAYNE.JOHNSON in English and Chinese KGs. Without the visual component, the two embedded vectors are far apart (top) while similar images can pull them together (bottom).

A major bottleneck for training EA models is the scarce cross-graph pivots² available as alignment signals (Chen et al. 2017; Sun et al. 2018). Besides, the sparsity of KGs is usually accompanied with weak structural correspondence, posing an even greater challenge to EA. To mitigate this problem, recent works have attempted to retrieve auxiliary supervision signals from the supplementary information of entities, such as attributes (Sun, Hu, and Li 2017; Trisedya, Qi, and Zhang 2019; Yang et al. 2020; Liu et al. 2020c) and descriptions (Chen et al. 2018). However, existing EA approaches are still limited in their capabilities. Our study proposes to leverage images, a natural component of entity profiles in many KGs (Lehmann et al. 2015; Vrandečić and Krötzsch 2014; Liu et al. 2019), for better EA. Images have been used to enrich entity representations for KG completion in a single-graph scenario (Xie et al. 2017; Mousselly-Sergieh et al. 2018; Peshkpour, Chen, and Singh 2018). However, the visual modality is yet to be explored for cross-graph tasks such as EA.

Our study stands upon several advantages that the visual

²In this paper, *pivot* is used interchangeably with *seed alignment* between cross-graph entities; *visual pivoting* means to use the visual space as intermediate to find seed alignment.

modality brings to EA. First, the visual concept of a named entity is usually universal, regardless of the language or the schema of the KG. Therefore, given a well-designed algorithm, images should provide the basis to find a set of reliable pivots. Second, images in KGs are freely-available and of high quality. Crucially, they are mostly manually verified and disambiguated. These abundant gold visual attributes in KGs render EA an ideal application scenario for visual representations. Third, images offer the possibility to enhance the representation of rare KG entities with impoverished structural contexts (Cao et al. 2020; Xiong et al. 2018; Hao et al. 2019). Images can be particularly beneficial in this setting, as entities of lower frequencies tend to be more concrete concepts (Hessel, Mimno, and Lee 2018) with stable visual representations (Kiela et al. 2014; Hewitt et al. 2018). To demonstrate the benefit from injecting images, we present a challenging example in Fig. 1. Without images, it is harder to infer the correspondence between DWAYNE_JOHNSON and its Chinese counterpart (“The Rock Johnson”) due to their dissimilar neighbourhoods in the two KGs. An alignment can be more easily induced by detecting visual similarity.

In this work, we propose **EVA** (Entity Visual Alignment), which incorporates images along with structures, relations and attributes to align entities in different KGs. During training, a learnable attention weighting scheme helps the alignment model to decide on the importance of each modality, and also provides interpretation for each modality’s contribution. As we show, an advantage of our approach is that the model is able to be trained on either a small set of seed alignment labels as in previous methods (semi-supervised setting), or using only a set of automatically induced visual pivots (unsupervised setting). Iterative learning (IL) is applied to expand the set of training pivots under both settings. On two large-scale standard benchmarks, i.e. DBP15k for cross-lingual EA and DWY15k for monolingual EA, **EVA** variants with or without alignment labels consistently outperform competitive baseline approaches.

The contributions of this work are three-fold: (i) We conduct the first investigation into the use of images as part of entity representations for EA, and achieve state-of-the-art (SOTA) performance across all settings. (ii) We leverage visual similarities to propose a fully unsupervised EA setting, avoiding reliance on any gold labels. Our model under the unsupervised setting performs closely to its semi-supervised results, even surpassing the previous best semi-supervised methods. (iii) We offer interpretability in our study by conducting ablation studies on the contributions from each modality and a thorough error analysis. We also provide insights on images’ particular impact on long-tail KG entities.

2 Related Work

Our work is connected to two research topics. Each has a large body of work which we can only provide as a highly selected summary.

Entity alignment. While early work employed symbolic or schematic methods to address the EA problem (Wijaya, Talukdar, and Mitchell 2013; Suchanek, Abiteboul, and Senel-

art 2011), more recent attention has been paid to embedding-based methods. A typical method of such is MTRANSE (Chen et al. 2017), which jointly trains a translational embedding model (Bordes et al. 2013) to encode language-specific KGs in separate embedding spaces, and a transformation to align counterpart entities across embeddings. Following this methodology, later works span the following three lines of studies to improve on this task. The first is to use alternatives of embedding learning techniques. Those include more advanced relational models such as contextual translations (Sun et al. 2019), residual recurrent networks (RSN, Guo, Sun, and Hu 2019) and relational reflection transformation (Mao et al. 2020), as well as variants of graph neural networks (GNNs) such as GCN (Wang et al. 2018; Yang et al. 2019; Wu et al. 2019a,b), GAT (Sun et al. 2020a; Zhu et al. 2019) and multi-channel GNNs (Cao et al. 2019). The second line of research focuses on capturing the alignment of entities with limited labels, therefore incorporating semi-supervised or metric learning techniques such as bootstrapping (Sun et al. 2018), co-training (Chen et al. 2018; Yang et al. 2020) and optimal transport (Pei, Yu, and Zhang 2019). Besides, to compensate for limited supervision signals in alignment learning, another line of recent works retrieves auxiliary supervision from side information of entities. Such information include numerical attributes (Sun, Hu, and Li 2017; Trisedya, Qi, and Zhang 2019), literals (Zhang et al. 2019; Otani et al. 2018) and descriptions of entities (Chen et al. 2018; Gesese et al. 2019). A recent survey by Sun et al. (2020b) has systematically summarised works in these lines.

The main contribution of this paper is relevant to the last line of research. To the best of our knowledge, this is the first attempt to incorporate the visual modality for EA in KGs. It also presents an effective unsupervised solution to this task, without the need of alignment labels that are typically required in previous works.

Multi-modal KG embeddings. While incorporating perceptual qualities has been a hot topic for language representation learning for many years, few attempts have been made towards building multi-modal KG embeddings. Xie et al. (2017) and Thoma, Rettinger, and Both (2017) are among the first to incorporate translational KG embedding methods (Bordes et al. 2013) with external visual information. However, they mostly explore the joint embeddings on intrinsic tasks like word similarity and link prediction. Mousselly-Sergieh et al. (2018) improve the model of Xie et al. to incorporate both visual and linguistic features under a unified translational embedding framework. Pezeshkpour, Chen, and Singh (2018) and Oñoro-Rubio et al. (2019) also model the interplay of images and KGs. However, Pezeshkpour, Chen, and Singh focus specifically on KG completion. Oñoro-Rubio et al. treat images as first class citizens for tasks like answering vision-relational queries instead of building joint representation for images and entities. The aforementioned works all focus on single KG scenarios. As far as we know, we are the first to use the intermediate visual space for EA between KGs.

Note that in the context of embedding alignment, many studies have incorporated images in lexical or sentential rep-

representations to solve cross-lingual tasks such as bilingual lexicon induction (Vulić et al. 2016; Rotman, Vulić, and Reichart 2018; Sigurdsson et al. 2020) or cross-modal matching tasks such as text-image retrieval (Gella et al. 2017; Kiros, Chan, and Hinton 2018; Kiela, Wang, and Cho 2018). Beyond embedding alignment, the idea of visual pivoting is also popular in the downstream task of machine translation (Caglayan et al. 2016; Huang et al. 2016; Hitschler, Schamoni, and Riezler 2016; Specia et al. 2016; Calixto and Liu 2017; Barrault et al. 2018; Su et al. 2019), but it is beyond the scope of this study. All of these works are not designed to deal with relational data that are crucial to performing EA.

3 Method

We start describing our method by formulating the learning resources. A KG (\mathcal{G}) can be viewed as a set of triplets that are constructed with an entity vocabulary (\mathcal{E}) and a relation vocabulary (\mathcal{R}), i.e. $\mathcal{G} = \{(e_1, r, e_2) : r \in \mathcal{R}; e_1, e_2 \in \mathcal{E}\}$ where a triplet records the relation r between the head and tail entities e_1, e_2 . Let $\mathcal{G}_s = \mathcal{E}_s \times \mathcal{R}_s \times \mathcal{E}_s, \mathcal{G}_t = \mathcal{E}_t \times \mathcal{R}_t \times \mathcal{E}_t$ denote two individual KGs (to be aligned). Given a pair of entities $e_s \in \mathcal{E}_s$ from source KG and $e_t \in \mathcal{E}_t$ from target KG, the goal of EA is to learn a function $f(\cdot, \cdot; \theta) : \mathcal{E}_s \times \mathcal{E}_t \rightarrow \mathbb{R}$ parameterised by θ that can estimate the similarity of e_s and e_t . $f(e_s, e_t; \theta)$ should be high if e_s, e_t are describing the same identity and low if they are not. Note that \mathcal{E}_s and \mathcal{E}_t ensure 1-to-1 alignment (Chen et al. 2018; Sun et al. 2018), as to be congruent to the design of mainstream KBs (Lehmann et al. 2015; Rebele et al. 2016) where disambiguation of entities is granted. To build joint representation for entities, we consider auxiliary information including images, relations and attributes. Let \mathcal{I} denote the set of all images; $\mathbf{R} \in \mathbb{R}^{N \times d_R}, \mathbf{A} \in \mathbb{R}^{N \times d_A}$ denote the matrices of relation and attribute features.

To tackle the EA task, our method jointly conducts two learning processes. A multi-modal embedding learning process aims at encoding both KGs \mathcal{G}_s and \mathcal{G}_t in a shared embedding space. Each entity in the embedding space is characterised based on both the KG structures and auxiliary information including images. In the shared space, the alignment learning process seeks to precisely capture the correspondence between counterpart entities by Neighbourhood Component Analysis (NCA, Goldberger et al. 2005; Liu et al. 2020b) and iterative learning. Crucially, the alignment learning process can be unsupervised, i.e. pivots are automatically inferred from the visual representations of entities without the need of EA labels. The rest of this section introduces the technical details of both learning processes.

3.1 Multi-Modal KG Embeddings

Given entities from two KGs \mathcal{G}_s and \mathcal{G}_t , and the auxiliary data $\mathcal{I}, \mathbf{R}, \mathbf{A}$, this section details how they are embedded into low-dimensional vectors.

Graph structure embedding. To model the structural similarity of \mathcal{G}_s and \mathcal{G}_t , capturing both entity and relation proximity, we use Graph Convolutional Network (GCN) proposed

by Kipf and Welling (2017). Formally, a multi-layer GCN’s operation on the l -th layer can be formulated as:

$$\mathbf{H}^{(l+1)} = [\tilde{\mathbf{D}}^{-\frac{1}{2}} \tilde{\mathbf{M}} \tilde{\mathbf{D}}^{-\frac{1}{2}} \mathbf{H}^{(l)} \mathbf{W}^{(l)}]_+, \quad (1)$$

where $[\cdot]_+$ is the ReLU activation; $\tilde{\mathbf{M}} = \mathbf{M} + \mathbf{I}_N$ is the adjacency matrix of $\mathcal{G}_s \cup \mathcal{G}_t$ plus an identity matrix (self-connection); $\tilde{\mathbf{D}}$ is a trainable layer-specific weight matrix; $\mathbf{H}^{(l)} \in \mathbb{R}^{N \times D}$ is the output of the previous GCN layer where N is number of entities and D is the feature dimension; $\mathbf{H}^{(0)}$ is randomly initialised. We use the output of the last GCN layer as the graph structure embedding \mathbf{F}_G .

Visual embedding. We use RESNET-152 (He et al. 2016), pre-trained on the ImageNet (Deng et al. 2009) recognition task, as the feature extractor for all images. For each image, we do a forward pass and take the last layer’s output before logits as the image representation (the RESNET itself is not fine-tuned). The feature is sent through a trainable feed-forward layer for the final image embedding:

$$\mathbf{F}_I = \mathbf{W}_I \cdot \text{RESNET}(\mathcal{I}) + \mathbf{b}_I. \quad (2)$$

The CNN-extracted visual representation is expected to capture both low-level similarity and high-level semantic relatedness between images (Kiela and Bottou 2014).³

Relation and attribute embeddings. Yang et al. (2019) showed that modelling relations and attributes with GCNs could pollute entity representations due to noise from neighbours. Following their investigation, we adopt a simple feed-forward network for mapping relation and attribute features into low-dimensional spaces:

$$\begin{aligned} \mathbf{F}_R &= \mathbf{W}_R \cdot \mathbf{R} + \mathbf{b}_R; \\ \mathbf{F}_A &= \mathbf{W}_A \cdot \mathbf{A} + \mathbf{b}_A. \end{aligned} \quad (3)$$

Modality fusion. We first l_2 -normalise each feature matrix by row and then fuse multi-modal features by trainable weighted concatenation:

$$\mathbf{F}_J = \bigoplus_{i=1}^n \left[\frac{e^{w_i}}{\sum_{j=1}^n e^{w_j}} \cdot \mathbf{F}_i \right], \quad (4)$$

where n is the number of modalities; w_i is an attention weight for the i -th modality. They are sent to a softmax before being multiplied to each modality’s l_2 -normalised representation, ensuring that the normalised weights sum to 1.

3.2 Alignment Learning

On top of the multi-modal embeddings \mathbf{F}_J for all entities, we compute the similarity of all bi-graph entity pairs and align them using an NCA loss. The training set is expanded using iterative learning.

³We compared several popular pre-trained visual encoders but found no substantial difference (§4.3).

Embedding similarity. Let $\mathbf{F}_J^s, \mathbf{F}_J^t$ denote embeddings of the source and target entities \mathcal{E}_s and \mathcal{E}_t respectively. We compute their cosine similarity matrix $\mathbf{S} = \langle \mathbf{F}_J^s, \mathbf{F}_J^t \rangle \in \mathbb{R}^{|\mathcal{E}_s| \times |\mathcal{E}_t|}$, where each entry \mathbf{S}_{ij} corresponds to the cosine similarity between the i -th entity in \mathcal{E}_s and the j -th in \mathcal{E}_t .

NCA loss. Inspired by the NCA-based text-image matching approach proposed by Liu et al. (2020b), we adopt an NCA loss of a similar form. It uses both local and global statistics to measure importance of samples and punishes hard negatives with a soft weighting scheme. This seeks to mitigate the *hubness problem* (Radovanović, Nanopoulos, and Ivanović 2010) in an embedding space. The loss is formulated below:

$$\mathcal{L} = \frac{1}{N} \sum_{i=1}^N \left(\frac{1}{\alpha} \log \left(1 + \sum_{m \neq i} e^{\alpha \mathbf{S}_{mi}} \right) + \frac{1}{\alpha} \log \left(1 + \sum_{n \neq i} e^{\alpha \mathbf{S}_{in}} \right) - \log \left(1 + \beta \mathbf{S}_{ii} \right) \right), \quad (5)$$

where α, β are temperature scales; N is the number of pivots within the mini-batch. We apply such loss on each modality separately and also on the merged multi-modal representation as specified in eq. (4). The joint loss is written as:

$$\mathcal{L}_{\text{Joint}} = \sum_i^n \mathcal{L}_i + \mathcal{L}_{\text{Multi-modal}} \quad (6)$$

where \mathcal{L}_i represents the loss term for aligning the i -th modality; $\mathcal{L}_{\text{Multi-modal}}$ is applied on the merged representation \mathbf{F}_J and is used for training the modality weights only. The reason for having separate terms for different modalities is that we use different hyper-parameters to accommodate their drastically distinct feature distributions. For all terms we used $\beta = 10$, but we picked different α 's: $\alpha = 5$ for \mathcal{L}_G ; $\alpha = 15$ for $\mathcal{L}_R, \mathcal{L}_A, \mathcal{L}_I, \mathcal{L}_J$.

Iterative learning. To improve learning with very few training pivots, we incorporate an iterative learning (IL) strategy to propose more pivots from unaligned entities. In contrast to previous work (Sun et al. 2018), we add a probation technique. In detail, for every K_e epochs, we make a new round of proposal. Each pair of cross-graph entities that are mutual nearest neighbours is proposed and added into a candidate list. If a proposed entity pair remains mutual nearest neighbours throughout K_s consecutive rounds (i.e. the probation phase), we permanently add it into the training set. Therefore, the candidate list refreshes every $K_e \cdot K_s$ epochs. In practice, we find that the probation technique has made the pivot discovery process more stable.

3.3 Unsupervised Visual Pivoting

Previous EA methods require annotated pivots that may not be widely available across KGs (Zhuang et al. 2017; Chen et al. 2017). Our method, however, can naturally extend to an unsupervised setting where visual similarities are leveraged to

Algorithm 1: Visual pivot induction.

```

input : visual embeddings from entities in the two graphs
         ( $\mathbf{F}_I^1, \mathbf{F}_I^2$ ); pivot dictionary size ( $n$ )
output: pivot dictionary
1  $\mathbf{M} \leftarrow \langle \mathbf{F}_I^1, \mathbf{F}_I^2 \rangle$  ▷ get similarity matrix
2  $\mathbf{m}_s \leftarrow \text{sort}(\mathbf{M})$  ▷ sort elements of  $\mathbf{M}$ 
3  $\mathcal{S} \leftarrow \{\}$  ▷ initialise seed dictionary
4  $\mathcal{R}_u \leftarrow \{\}; \mathcal{C}_u \leftarrow \{\}$  ▷ for recording used row/column
5 while  $|\mathcal{S}| = n$  do
6    $m \leftarrow \mathbf{m}_s.\text{pop}()$  ▷ get the highest ranked score
7   if  $m.\text{ri} \notin \mathcal{R}_u \ \& \ m.\text{ci} \notin \mathcal{C}_u$  then
8      $\mathcal{S} \leftarrow \mathcal{S} \cup (m.\text{ri}, m.\text{ci})$  ▷ store the pair
9      $\mathcal{R}_u \leftarrow \mathcal{R}_u \cup m.\text{ci}$ 
10     $\mathcal{C}_u \leftarrow \mathcal{C}_u \cup m.\text{ri}$ 
11  end
12 end
13 return  $\mathcal{S}$  ▷ return the obtained visual pivot dictionary

```

infer correspondence between KGs, and no annotated cross-graph pivots are required. All cross-graph supervision comes from an automatically induced visual dictionary (visual pivots) containing the most visually alike cross-graph entities. Specifically, we first compute cosine similarities of all cross-graph entities' visual representations in the data set. Then we sort the cosine similarity matrix from high to low. We collect visual pivots starting from the most similar pairs. Once a pair of entities is collected, all other links associated with the two entities are discarded. In the end, we obtain a cross-graph pivot list that records the top- k visually similar entity pairs without repetition of entities. From these visual pivots, we apply iterative learning (§3.2) to expand the training set. The algorithm of obtaining visual pivots is formally described in Algorithm 1. Let n be the number of entities in one language, the algorithm takes $\mathcal{O}(n^2 \log(n^2) + n^2) = \mathcal{O}(n^2 \log n)$.

Our approach is related to some recent efforts on word translation with images (Bergsma and Van Durme 2011; Kiela, Vulić, and Clark 2015; Hewitt et al. 2018). However, those efforts focus on obtaining cross-lingual parallel signals from web-crawled images provided by search engines (e.g. Google Image Search). This can result in noisy data caused by issues like ambiguity in text. For example, for a query *mouse*, the search engine might return images for both the rodent and the computer mouse. Visual pivoting is thus more suitable in the context of EA, as images provided by KGs are mostly human-verified and disambiguated, serving as gold visual representations of the entities. Moreover, cross-graph entities are not necessarily cross-lingual, meaning that the technique could also benefit a monolingual scenario.

4 Experiments

In this section, we conduct experiments on two benchmark data sets (§4.1), under both semi- and unsupervised settings (§4.2). We also provide detailed ablation studies on different model components (§4.3), and study the impact of incorporating visual representations on long-tail entities (§4.4).

model	FR→EN			JA→EN			ZH→EN		
	H@1	H@10	MRR	H@1	H@10	MRR	H@1	H@10	MRR
MTRANSE (Chen et al. 2017)	.224	.556	.335	.279	.575	.349	.308	.614	.364
JAPE (Sun, Hu, and Li 2017)	.324	.667	.430	.363	.685	.476	.412	.745	.490
GCN (Wang et al. 2018)	.373	.745	.532	.399	.745	.546	.413	.744	.549
MUGNN (Cao et al. 2019)	.495	.870	.621	.501	.857	.621	.494	.844	.611
RSN (Guo, Sun, and Hu 2019)	.516	.768	.605	.507	.737	.590	.508	.745	.591
KECG (Li et al. 2019)	.486	.851	.610	.490	.844	.610	.478	.835	.598
HMAN (Yang et al. 2019)	.543	.867	-	.565	.866	-	.537	.834	-
GCN-JE (Wu et al. 2019b)	.483	.778	-	.466	.746	-	.459	.729	-
GMN (Xu et al. 2019)*	.596	.876	.679	.465	.728	.580	.433	.681	.479
ALINET (Sun et al. 2020a)	.552	.852	.657	.549	.831	.645	.539	.826	.628
EVA w/o IL	.715	.936	.795	.716	.926	.792	.720	.925	.793
	±.003	±.002	±.004	±.008	±.004	±.006	±.004	±.006	±.003
BOOTEA (Sun et al. 2018)	.653	.874	.731	.622	.854	.701	.629	.847	.703
MMEA (Shi and Xiao 2019)	.635	.878	-	.623	.847	-	.647	.858	-
NAEA (Zhu et al. 2019)	.673	.894	.752	.641	.873	.718	.650	.867	.720
EVA w/ IL	.793	.942	.847	.762	.913	.817	.761	.907	.814
	±.005	±.002	±.004	±.008	±.003	±.006	±.008	±.005	±.006

Table 1: Cross-lingual EA results on DBP15k. Comparison with related works with and without using IL. “-” means not reported by the original paper. “*” indicates our reproduced results for which any use of machine translation or cross-lingual alignment labels other than those provided in the benchmark are removed.




Figure 2: Unsupervised EVA vs. semi-supervised EVA. Plotting H@1 against number of induced visual seeds for jump-starting the training.

4.1 Experimental Settings

Data sets. The experiments are conducted on DBP15k (Sun, Hu, and Li 2017) and DWY15k (Guo, Sun, and Hu 2019). DBP15k is a widely used cross-lingual EA benchmark. It contains four language-specific KGs from DBpedia, and has three bilingual EA settings, i.e., French-English (FR-EN), Japanese-English (JA-EN) and Chinese-English (ZH-EN). DBpedia has also released images for English, French and Japanese versions. Note that since Chinese images are not released in DBpedia, we extracted them from the raw Chinese Wikipedia dump with the same process as described by Lehmann et al. (2015). DWY15k is a monolingual data set, focusing on EA for DBpedia-Wikidata and DBpedia-YAGO. It has two subsets, DWY15k-norm and DWY15k-dense, whereof the former is much sparser. As YAGO does not have image components, we experiment on DBpedia-Wikidata only. Note that not all but ca. 50-85% entities have images, as shown in §4.4. For an entity without an image, we assign a random vector sampled from a normal distribution, parameterised by the mean and standard deviation of other images. As for relation and attribute features, we extract them in the same way as Yang et al. (2019).

Model configurations. The GCN has two layers with input, hidden and output dimensions of 400, 400, 200 respectively. Attribute and relation features are mapped to 100- d . Images are transformed to 2048- d features by RESNET and then mapped to 200- d . For model variants without IL, training is limited to 500 epochs. Otherwise, after the first 500 epochs, IL is conducted for another 500 epochs with the configurations $K_e = 5$, $K_s = 10$ as described in §3.2. We train all models using a batch size of 7,500. The models are optimised using AdamW (Loshchilov and Hutter 2019) with a learning rate of $5e-4$ and a weight decay of $1e-2$. More implementation details are available in the Appendix A of (Liu et al. 2020a).

Evaluation protocols. Following convention, we report three metrics on both data sets, including $H@\{1,10\}$ (the proportion of ground truth being ranked no further than top $\{1,10\}$), and MRR (mean reciprocal rank). During inference, we use Cross-domain Similarity Local Scaling (CSLS; Lample et al. 2018) to post-process the cosine similarity matrix, which is employed by default in some recent works (Sun et al. 2019, 2020a). $k = 3$ is used for defining local neighbourhood of CSLS. All models are run for 5 times with 5 different random seeds and the average with variances are reported. **Bold**

numbers in tables come from the best models and **underline** means with statistical significance (p-value < 0.05 in t-test).

The baseline results on DBP15k come from ten methods with IL, and three without IL. We accordingly report the results by **EVA** with and without IL. Note that a few methods may incorporate extra cross-lingual alignment labels by initialising training with machine translation (Wu et al. 2019b; Yang et al. 2019) or pre-aligned word vectors (Xu et al. 2019). For fair comparison in this study, we report results from the versions of these methods that do not use any alignment signals apart from the training data. On DWY15k, there are also two settings in the literature, different in whether to use the surface form embeddings of monolingual entities (Yang et al. 2020) or not (Guo, Sun, and Hu 2019). We report results from **EVA** with and without using surface forms⁴, and compare with five SOTA baselines.

4.2 Main Results

Semi-supervised EA. §3.3 reports the results on semi-supervised cross-lingual EA. This setting compares **EVA** with baseline methods using the original data split of DBP15k, i.e., using 30% of the EA labels for training. Consequently, **EVA** achieves SOTA performance and surpasses baseline models drastically both with or without IL. Specifically, in the w/o IL setting, **EVA** leads to 12.3-17.6% absolute improvement in H@1 over the best baseline. When incorporating IL, **EVA** gains 11.9-12.5% absolute improvement in H@1 over the best IL-based baseline method. This indicates that incorporating the visual representation competently improves the cross-lingual entity representations for inferring their correspondences, without the need of additional supervision labels. The results on monolingual EA generally exhibit similar observations. As reported in §4.2, without incorporating surface form information, **EVA** surpasses the strongest baseline method by 16.8% in H@1 on the normal split and 4.2% on the dense split. With surface forms considered, **EVA** offers near-perfect results, outperforming the SOTA method by 26.8% in H@1 on the normal split and 6.6% in H@1 on the dense split. The experiments here indicate that, **EVA** is able to substantially improve SOTA EA systems under both monolingual and cross-lingual settings.

Unsupervised EA. We also use the visual pivoting technique (§3.3) with **EVA** to conduct unsupervised EA, without using any annotated alignment labels. We compare **EVA**’s best unsupervised and semi-supervised results in §4.2. The unsupervised **EVA** yields 1.9-6.3% lower H@1 than the semi-supervised version, but still notably outperforms the best semi-supervised baseline (§3.3) by 5.6-10.1%. We change the number of visual seeds by adjusting the threshold n in Algorithm 1, and test the model’s sensitiveness to the

⁴When incorporating surface form, we use FASTTEXT (Bojanowski et al. 2017) to embed surface strings into low-dimensional vectors $\mathbf{S} \in \mathbb{R}^{N \times d_s}$ ($d_s = 300$), and learn a linear transformation to obtain final representations in 100-d: $\mathbf{F}_S = \mathbf{W}_S \cdot \mathbf{S} + \mathbf{b}_S$. We merge and train the surface form modality in the same way as the other modalities.

model	DBP→WD (N)			DBP→WD (D)		
	H@1	H@10	MRR	H@1	H@10	MRR
BOOTEA (Sun et al. 2018)	.323	.631	.420	.678	.912	.760
GCN (Wang et al. 2018)	.177	.378	.250	.431	.713	.530
W/O SF JAPE (Sun, Hu, and Li 2017)	.219	.501	.310	.393	.705	.500
W/O SF RSN (Guo, Sun, and Hu 2019)	.388	.657	.490	.763	.924	.830
W/O SF COTSAE (Yang et al. 2020)	.423	.703	.510	.823	.954	.870
EVA w/o SF	.593	.775	.655	.874	.962	.908
	±.004 ±.005 ±.003			±.002 ±.003 ±.002		
W/SF COTSAE (Yang et al. 2020)	.709	.904	.770	.922	.983	.940
W/SF EVA w/ SF	.985	.995	.989	.994	1.0	.996
	±.001 ±.000 ±.001			±.001 ±.001 ±.000		

Table 2: Monolingual EA results on DWY15k-DW (N: normal split; D: dense split). **EVA** using IL is compared with related works with and without using surface forms (W/ SF & W/O SF).

setting	FR→EN			JA→EN			ZH→EN		
	H@1	H@10	MRR	H@1	H@10	MRR	H@1	H@10	MRR
Unsup.	.731	.909	.792	.737	.890	.791	.752	.895	.804
	±.004 ±.003 ±.003			±.008 ±.004 ±.006			±.006 ±.004 ±.005		
Semi-sup.	.793	.942	.847	.762	.913	.817	.761	.907	.814
	±.003 ±.002 ±.004			±.008 ±.004 ±.006			±.004 ±.006 ±.003		

Table 3: Comparing unsupervised and semi-supervised **EVA** results on DBP15k.

threshold. As shown in Fig. 2, the optimal seed size is 4k on FR→EN and 6k on JA→EN and ZH→EN.

It is worth noticing that a good alignment (H@1>55%) can be obtained using as few as a hundred visual seeds. As the number of seeds grows, the model gradually improves, reaching >70% H@1 with more than 3k seeds. Then the scores plateau for a period and start to decrease with more than 4k (on FR→EN) or 6k (JA→EN, ZH→EN) seeds. This is because a large visual seed dictionary starts to introduce noise. Empirically, we find that a 0.85 cosine similarity threshold is a good cut-off point.

4.3 Ablation Study

We report an ablation study of **EVA** in §4.3 using DBP15k (FR→EN). As shown, IL brings ca. 8% absolute improvement. This gap is smaller than what has been reported previously (Sun et al. 2018). This is because the extra visual supervision in our method already allows the model to capture fairly good alignment in the first 500 epochs, leaving smaller room for further improvement from IL. CSLS gives minor but consistent improvement to all metrics during inference. While CSLS is mainly used to reduce hubs in a dense space such as textual embeddings (Lample et al. 2018), we suspect that it cannot bring substantial improvement to our sparse multi-modal space. Besides, the hubness problem is already partly tackled by our NCA loss. The sparseness of multi-modal space can also explain our choice of $k = 3$, which we found to be better than the previous $k = 10$.

Regarding the impact from different modalities, structure

model	H@1	H@10	MRR
W/O structure	.391 ±.004	.514 ±.003	.423 ±.004
W/O image	.749 ±.002	.929 ±.002	.817 ±.001
W/O attribute	.750 ±.003	.927 ±.001	.813 ±.003
W/O relation	.763 ±.006	.928 ±.003	.823 ±.004
W/O IL	.715 ±.003	.936 ±.002	.795 ±.004
W/O CSLS	.786 ±.005	.928 ±.001	.838 ±.003
full model	.793 ±.003	.942 ±.002	.847 ±.004

Table 4: Ablation study of **EVA** based on DBP15k (FR→EN).

visual encoder	H@1	H@10	MRR
ResNet50	.713 ±.003	.938 ±.002	.794 ±.003
ResNet50 (places365)	.710 ±.002	.937 ±.002	.792 ±.002
ResNet152	.715 ±.003	.936 ±.002	.795 ±.004
DenseNet201	.716 ±.005	.935 ±.002	.796 ±.003
Inception v3	.711 ±.002	.936 ±.002	.792 ±.002

Table 5: Comparing visual encoders. Results are obtained on DBP15k (FR→EN) using **EVA** w/o IL.

remains the most important for our model. Dropping structural embedding decreases H@1 from ca. 80% to below 40%, cutting the performance by half. This is in line with the findings by Yang et al. (2019). Image, attributes and relations are of similar importance. The removal of images and attributes decrease H@1 by 4-5% while removing relations causes ca. 3% drop in H@1. This general pattern roughly corresponds to the modality attention weights. On DBP15k, while all weights start at 0.25, after training, they become ca. 0.45, 0.21, 0.17 and 0.16 for structures, images, relations and attributes respectively.

In addition, we explore several other popular options of pre-trained visual encoder architectures including ResNet (He et al. 2016), GoogLeNet (Szegedy et al. 2015), DenseNet (Huang et al. 2017) and Inception v3 (Szegedy et al. 2016) as feature extractors for images. One of the variants, ResNet (Places365) (Zhou et al. 2017), is pre-trained on a data set from the outdoor-scene domain and is expected to be better at capturing location-related information. In general, we found little difference in model performance with different visual encoders. As suggested in §4.3, variances from different models across different metrics are generally < 1%. All numbers reported in the paper have been using ResNet152 as it is one of the most widely used visual feature extractor. It is also possible to fine-tune the visual encoder with the full model in an end-to-end fashion. However, the computation cost would be extremely large under our setting.

4.4 Analysis on Long-tail Entities

Like lexemes in natural languages, the occurrence of entities in KG triplets also follow a long-tailed distribution (Fig. 3). Long-tail entities are poorly connected to others in the graph and thus have less structural information for inducing reliable representation and alignment. We argue that images might remedy the issue by providing alternative sources of signal for representing these long-tail entities. To




Figure 3: Long-tailed distribution of entity appearances in KG triplets, using 100 randomly sampled entities in DBP15k (FR-EN).

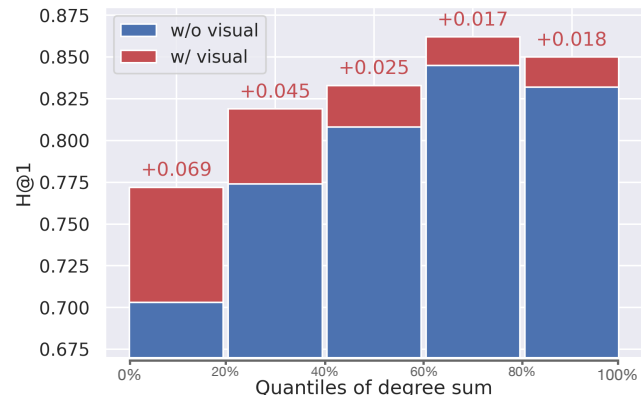


Figure 4: Plotting H@1 against different test splits on FR-EN (frequency low to high from left to right). Models w/ or w/o visual information are compared. The plot suggests that images have improved long-tail entities’ alignment more.


validate our hypothesis, we stratify the test set of DBP15k (FR→EN) into five splits of entity pairs based on their degree centrality in the graphs. Specifically, for all entity pairs $(e_s, e_t) \in \mathcal{G}_s \times \mathcal{G}_t$ in the test set, we sort them by their degree sum, i.e., $\text{DegSum}(e_s, e_t) := \text{deg}(e_s) + \text{deg}(e_t)$, and split them into five sets of equal sizes, corresponding to five ranges of DegSum partitioned by 14, 18, 23 and 32, respectively. Across the five splits, we compare the performance by **EVA** (w/o IL) against its variant where visual inputs are disabled. The results in Fig. 4 suggest that entities in the lower ranges of degree centrality benefit more from the visual representations. This demonstrates that the visual modality particularly enhances the match of long-tail entities which gain less information from other modalities.

As an example, in DBP15k (FR→EN), the long-tail entity *Stade_olympique_de_Munich* has only three occurrences in French. The top three retrieved entities in English by **EVA** w/o visual representation are *Olympic_Stadium_(Amsterdam)*, *Friends_Arena* and *Olympiastadion_(Munich)*. The embedding without visual information was only able to narrow down the answer to European stadiums, but failed to correctly order the specific stadiums (Fig. 5). With the visual cues, **EVA** is able to rank the correct item as the top 1.


Note that in Fig. 4, the split of the most frequent entities (80-100% quantiles) generally displays worse performance

	FR↔EN		JA↔EN		ZH↔EN		DBP↔WD (norm)		DBP↔WD (dense)	
	FR	EN	JA	EN	ZH	EN	DBP	WD	DBP	WD
image covered	14,174	13,858	12,739	13,741	15,912	14,125	8,517	8,791	7,744	7,315
all entities	19,661	19,993	19,814	19,780	19,388	19,572	15,000	15,000	15,000	15,000

Table 6: Image coverage statistics on DBP15k and DWY15k. The image coverage of DBP15k (ca. 65-85%) is generally better than DWY15k (ca. 50-60%).



(a) Query (French): *Stade_olympique_de_Munich*



(b)

(c)

(d)

Figure 5: **EVA** w/o images ranks (b) *Olympic_Stadium_(Amsterdam)* at top 1, (c) *Friends_Arena*, (d) *Olympiastadion_(Munich)* at top 2, 3 respectively. Through visual disambiguation, **EVA** ranks the correct concept (d) at top 1.

than the second most frequent split (60-80% quantiles), suggesting that, a denser neighbourhood does not always lead to better alignment. This is consistent with Sun et al. (2020a)’s observation that, entities with high degree centrality may be affected by the heterogeneity of their neighbourhood.

4.5 Error Analysis

On the monolingual setting (DBP15k-WD), **EVA** has reached near-perfect performance. On the cross-lingual setting (DBP15k), however, there is still a $>20\%$ gap from perfect alignment. One might wonder why the involvement of images has not solved the remaining 20% errors. By looking into the errors made by **EVA** (w/o IL) on DBP15k (FR→EN), we observe that among the 2955 errors, 1945 (i.e., ca. 2/3) of them are entities without valid images. In fact, only 50-70% of entities in our study have images, according to §4.4. This is inherent to knowledge bases themselves and cannot be easily resolved without an extra step of linking the entities to some external image database. For the remaining 1k errors, ca. 40% were wrongly predicted regardless of with or without images. The other 60% were correctly predicted before injecting visual information, but were missed when images were present. Such errors can be mainly attributed to the consistency/robustness issues in visual representations especially for more abstract entities as they tend to have multiple plausible visual representations. Here is a real example: *Université_de_Varsovie* (French) has the photo of its front gate in

the profile while its English equivalent *University_of_Warsaw* uses its logo in the profile. The drastically different visual representations cause a misalignment. While images are in most cases helpful for the alignment, it requires further investigation for a mechanism to filter out the small fraction of unstable visual representations. This is another substantial research direction for future work.

5 Conclusion

We propose a new model **EVA** that uses images as pivots for aligning entities in different KGs. Through an attention-based modality weighting scheme, we fuse multi-modal information from KGs into a joint embedding and allow the alignment model to automatically adjust modality weights. Besides experimenting with the traditional semi-supervised setting, we present an unsupervised approach, where **EVA** leverages visual similarities of entities to build a seed dictionary from scratch and expand the dictionary with iterative learning. The semi-supervised **EVA** claims new SOTA on two EA benchmarks, surpassing previous methods by large margins. The unsupervised **EVA** achieves $>70\%$ accuracy, being close to its performance under the semi-supervised setting, and outperforming the previous best semi-supervised baseline. Finally, we conduct thorough ablation studies and error analysis, offering insights on the benefits of incorporating images for long-tail KG entities. The implication of our work is that perception is a crucial element in learning entity representation and associating knowledge. In this way, our work also highlights the necessity of fusing different modalities in developing intelligent learning systems (Mooney 2008).

Acknowledgments

We appreciate the anonymous reviewers for their insightful comments and suggestions. Fangyu Liu is supported by Grace & Thomas C.H. Chan Cambridge Scholarship. This research is supported in part by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via IARPA Contract No. 2019-19051600006 under the BETTER Program, and by Contracts HR0011-18-2-0052 and FA8750-19-2-1004 with the US Defense Advanced Research Projects Agency (DARPA). The views expressed are those of the authors and do not reflect the official policy or position of the Department of Defense or the U.S. Government.

References

Barraut, L.; Bougares, F.; Specia, L.; Lala, C.; Elliott, D.; and Frank, S. 2018. Findings of the Third Shared Task on Multimodal Machine Translation. In *WMT*, 304–323.

- Bergsma, S.; and Van Durme, B. 2011. Learning bilingual lexicons using the visual similarity of labeled web images. In *IJCAI*.
- Bleiholder, J.; and Naumann, F. 2009. Data fusion. *ACM Computing Surveys* 41(1): 1–41.
- Bojanowski, P.; Grave, E.; Joulin, A.; and Mikolov, T. 2017. Enriching Word Vectors with Subword Information. *TACL* 5: 135–146.
- Bollacker, K.; Evans, C.; Paritosh, P.; Sturge, T.; and Taylor, J. 2008. Freebase: a collaboratively created graph database for structuring human knowledge. In *SIGMOD*, 1247–1250.
- Bordes, A.; Usunier, N.; Garcia-Duran, A.; Weston, J.; and Yakhnenko, O. 2013. Translating embeddings for modeling multi-relational data. In *NeurIPS*, 2787–2795.
- Bryl, V.; and Bizer, C. 2014. Learning conflict resolution strategies for cross-language wikipedia data fusion. In *WWW*, 1129–1134.
- Caglayan, O.; Aransa, W.; Wang, Y.; Masana, M.; García-Martínez, M.; Bougares, F.; Barrault, L.; and van de Weijer, J. 2016. Does Multimodality Help Human and Machine for Translation and Image Captioning? In *WMT*, 627–633.
- Calixto, I.; and Liu, Q. 2017. Incorporating Global Visual Features into Attention-based Neural Machine Translation. In *EMNLP*, 992–1003.
- Cao, E.; Wang, D.; Huang, J.; and Hu, W. 2020. Open Knowledge Enrichment for Long-tail Entities. In *WWW*, 384–394.
- Cao, Y.; Liu, Z.; Li, C.; Li, J.; and Chua, T.-S. 2019. Multi-Channel Graph Neural Network for Entity Alignment. In *ACL*, 1452–1461.
- Chen, M.; Tian, Y.; Chang, K.-W.; Skiena, S.; and Zaniolo, C. 2018. Co-training embeddings of knowledge graphs and entity descriptions for cross-lingual entity alignment. In *IJCAI*, 3998–4004.
- Chen, M.; Tian, Y.; Yang, M.; and Zaniolo, C. 2017. Multilingual knowledge graph embeddings for cross-lingual knowledge alignment. In *IJCAI*, 1511–1517.
- Cui, W.; Xiao, Y.; Wang, H.; Song, Y.; Hwang, S.-w.; and Wang, W. 2017. KBQA: learning question answering over QA corpora and knowledge bases. *PVLDB* 10(5): 565–576.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. ImageNet: A large-scale hierarchical image database. In *CVPR*, 248–255. IEEE.
- Gella, S.; Sennrich, R.; Keller, F.; and Lapata, M. 2017. Image Pivoting for Learning Multilingual Multimodal Representations. In *EMNLP*, 2839–2845.
- Gesese, G. A.; Biswas, R.; Alam, M.; and Sack, H. 2019. A Survey on Knowledge Graph Embeddings with Literals: Which model links better Literal-ly? *Semantic Web Journal*.
- Goldberger, J.; Hinton, G. E.; Roweis, S. T.; and Salakhutdinov, R. R. 2005. Neighbourhood components analysis. In *NeurIPS*, 513–520.
- Guo, L.; Sun, Z.; and Hu, W. 2019. Learning to Exploit Long-term Relational Dependencies in Knowledge Graphs. In *ICML*, 2505–2514.
- Hao, J.; Chen, M.; Yu, W.; Sun, Y.; and Wang, W. 2019. Universal representation learning of knowledge bases by jointly embedding instances and ontological concepts. In *KDD*, 1709–1719.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *CVPR*, 770–778.
- Hessel, J.; Mimno, D.; and Lee, L. 2018. Quantifying the Visual Concreteness of Words and Topics in Multimodal Datasets. In *NAACL*, 2194–2205.
- Hewitt, J.; Ippolito, D.; Callahan, B.; Kriz, R.; Wijaya, D. T.; and Callison-Burch, C. 2018. Learning translations via images with a massively multilingual image dataset. In *ACL*, 2566–2576.
- Hitschler, J.; Schamoni, S.; and Riezler, S. 2016. Multimodal Pivots for Image Caption Translation. In *ACL*, 2399–2409.
- Hoffmann, R.; Zhang, C.; Ling, X.; Zettlemoyer, L.; and Weld, D. S. 2011. Knowledge-based weak supervision for information extraction of overlapping relations. In *ACL-HLT*, 541–550.
- Huang, G.; Liu, Z.; Van Der Maaten, L.; and Weinberger, K. Q. 2017. Densely connected convolutional networks. In *CVPR*, 4700–4708.
- Huang, P.-Y.; Liu, F.; Shiang, S.-R.; Oh, J.; and Dyer, C. 2016. Attention-based multimodal neural machine translation. In *WMT*.
- Kiela, D.; and Bottou, L. 2014. Learning image embeddings using convolutional neural networks for improved multi-modal semantics. In *EMNLP*, 36–45.
- Kiela, D.; Hill, F.; Korhonen, A.; and Clark, S. 2014. Improving multi-modal representations using image dispersion: Why less is sometimes more. In *ACL*, 835–841.
- Kiela, D.; Vulić, I.; and Clark, S. 2015. Visual Bilingual Lexicon Induction with Transferred ConvNet Features. In *EMNLP*, 148–158.
- Kiela, D.; Wang, C.; and Cho, K. 2018. Dynamic Meta-Embeddings for Improved Sentence Representations. In *EMNLP*.
- Kipf, T. N.; and Welling, M. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *ICLR*.
- Kiros, J.; Chan, W.; and Hinton, G. 2018. Illustrative language understanding: Large-scale visual grounding with image search. In *ACL*, 922–933.
- Koncel-Kedziorski, R.; Bekal, D.; Luan, Y.; Lapata, M.; and Hajishirzi, H. 2019. Text Generation from Knowledge Graphs with Graph Transformers. In *NAACL*, 2284–2293.
- Lample, G.; Conneau, A.; Ranzato, M.; Denoyer, L.; and Jégou, H. 2018. Word translation without parallel data. In *ICLR*.
- Lehmann, J.; Isele, R.; Jakob, M.; Jentzsch, A.; Kontokostas, D.; Mendes, P. N.; Hellmann, S.; Morse, M.; Van Kleef, P.; Auer, S.; et al. 2015. DBpedia—a large-scale, multilingual knowledge base extracted from Wikipedia. *Semantic Web* 6(2): 167–195.
- Li, C.; Cao, Y.; Hou, L.; Shi, J.; Li, J.; and Chua, T.-S. 2019. Semi-supervised Entity Alignment via Joint Knowledge Embedding Model and Cross-graph Model. In *EMNLP-IJCNLP*, 2723–2732.
- Liu, F.; Chen, M.; Roth, D.; and Collier, N. 2020a. Visual Pivoting for (Unsupervised) Entity Alignment. *arXiv preprint arXiv:2009.13603*.
- Liu, F.; Ye, R.; Wang, X.; and Li, S. 2020b. HAL: Improved Text-Image Matching by Mitigating Visual Semantic Hubs. In *AAAI*.
- Liu, Y.; Li, H.; Garcia-Duran, A.; Niepert, M.; Onoro-Rubio, D.; and Rosenblum, D. S. 2019. MMKG: Multi-modal Knowledge Graphs. In *ESWC*, 459–474. Springer.
- Liu, Z.; Cao, Y.; Pan, L.; Li, J.; and Chua, T.-S. 2020c. Exploring and Evaluating Attributes, Values, and Structure for Entity Alignment. In *EMNLP*, 6355–6364.
- Loshchilov, I.; and Hutter, F. 2019. Decoupled Weight Decay Regularization. In *ICLR*.
- Mao, X.; Wang, W.; Xu, H.; Wu, Y.; and Lan, M. 2020. Relational Reflection Entity Alignment. In *CIKM*, 1095–1104.
- Mooney, R. J. 2008. Learning to Connect Language and Perception. In *AAAI*, 1598–1601.

- Mousselly-Sergieh, H.; Botschen, T.; Gurevych, I.; and Roth, S. 2018. A multimodal translation-based approach for knowledge graph representation learning. In **SEM*, 225–234.
- Oñoro-Rubio, D.; Niepert, M.; García-Durán, A.; González-Sánchez, R.; and López-Sastre, R. J. 2019. Answering Visual-Relational Queries in Web-Extracted Knowledge Graphs. In *AKBC*.
- Otani, N.; Kiyomaru, H.; Kawahara, D.; and Kurohashi, S. 2018. Cross-lingual Knowledge Projection Using Machine Translation and Target-side Knowledge Base Completion. In *COLING*, 1508–1520.
- Pei, S.; Yu, L.; and Zhang, X. 2019. Improving cross-lingual entity alignment via optimal transport. In *IJCAI*, 3231–3237. AAAI Press.
- Pezeshkpour, P.; Chen, L.; and Singh, S. 2018. Embedding Multi-modal Relational Data for Knowledge Base Completion. In *EMNLP*, 3208–3218.
- Radhakrishnan, P.; Talukdar, P.; and Varma, V. 2018. ELDEN: Improved entity linking using densified knowledge graphs. In *NAACL*.
- Radovanović, M.; Nanopoulos, A.; and Ivanović, M. 2010. Hubs in space: Popular nearest neighbors in high-dimensional data. *JMLR* 11(Sep): 2487–2531.
- Rebele, T.; Suchanek, F.; Hoffart, J.; Biega, J.; Kuzey, E.; and Weikum, G. 2016. YAGO: A multilingual knowledge base from wikipedia, wordnet, and geonames. In *ISWC*, 177–185. Springer.
- Rotman, G.; Vulić, I.; and Reichart, R. 2018. Bridging languages through images with deep partial canonical correlation analysis. In *ACL*, 910–921.
- Shi, X.; and Xiao, Y. 2019. Modeling Multi-mapping Relations for Precise Cross-lingual Entity Alignment. In *EMNLP-IJCNLP*, 813–822.
- Sigurðsson, G. A.; Alayrac, J.-B.; Nematzadeh, A.; Smaira, L.; Malinowski, M.; Carreira, J.; Blunsom, P.; and Zisserman, A. 2020. Visual Grounding in Video for Unsupervised Word Translation. *CVPR*.
- Specia, L.; Frank, S.; Sima'an, K.; and Elliott, D. 2016. A shared task on multimodal machine translation and crosslingual image description. In *WMT*, 543–553.
- Su, Y.; Fan, K.; Bach, N.; Kuo, C.-C. J.; and Huang, F. 2019. Unsupervised multi-modal neural machine translation. In *CVPR*.
- Suchanek, F. M.; Abiteboul, S.; and Senellart, P. 2011. PARIS: probabilistic alignment of relations, instances, and schema. *PVLDB* 5(3): 157–168.
- Sun, Z.; Hu, W.; and Li, C. 2017. Cross-lingual entity alignment via joint attribute-preserving embedding. In *ISWC*, 628–644. Springer.
- Sun, Z.; Hu, W.; Zhang, Q.; and Qu, Y. 2018. Bootstrapping Entity Alignment with Knowledge Graph Embedding. In *IJCAI*.
- Sun, Z.; Huang, J.; Hu, W.; Chen, M.; Guo, L.; and Qu, Y. 2019. TransEdge: Translating Relation-Contextualized Embeddings for Knowledge Graphs. In *ISWC*, 612–629. Springer.
- Sun, Z.; Wang, C.; Hu, W.; Chen, M.; Dai, J.; Zhang, W.; and Qu, Y. 2020a. Knowledge Graph Alignment Network with Gated Multi-hop Neighborhood Aggregation. In *AAAI*.
- Sun, Z.; Zhang, Q.; Hu, W.; Wang, C.; Chen, M.; Akrami, F.; and Li, C. 2020b. A Benchmarking Study of Embedding-based Entity Alignment for Knowledge Graphs. *PVLDB* 13.
- Szegedy, C.; Liu, W.; Jia, Y.; Sermanet, P.; Reed, S.; Anguelov, D.; Erhan, D.; Vanhoucke, V.; and Rabinovich, A. 2015. Going deeper with convolutions. In *CVPR*, 1–9.
- Szegedy, C.; Vanhoucke, V.; Ioffe, S.; Shlens, J.; and Wojna, Z. 2016. Rethinking the inception architecture for computer vision. In *CVPR*, 2818–2826.
- Thoma, S.; Rettinger, A.; and Both, F. 2017. Towards holistic concept representations: Embedding relational knowledge, visual attributes, and distributional word semantics. In *ISWC*, 694–710.
- Trisedya, B. D.; Qi, J.; and Zhang, R. 2019. Entity alignment between knowledge graphs using attribute embeddings. In *AAAI*, volume 33, 297–304.
- Vrandečić, D.; and Krötzsch, M. 2014. Wikidata: a free collaborative knowledgebase. *Communications of the ACM* 57(10): 78–85.
- Vulić, I.; Kiela, D.; Clark, S.; and Moens, M. F. 2016. Multi-modal representations for improved bilingual lexicon learning. In *ACL*, 188–194.
- Wang, Z.; Lv, Q.; Lan, X.; and Zhang, Y. 2018. Cross-lingual knowledge graph alignment via graph convolutional networks. In *EMNLP*, 349–357.
- Wijaya, D.; Talukdar, P. P.; and Mitchell, T. 2013. PIDGIN: ontology alignment using web text as interlingua. In *CIKM*, 589–598.
- Wu, Y.; Liu, X.; Feng, Y.; Wang, Z.; Yan, R.; and Zhao, D. 2019a. Relation-aware entity alignment for heterogeneous knowledge graphs. In *IJCAI*, 5278–5284. AAAI Press.
- Wu, Y.; Liu, X.; Feng, Y.; Wang, Z.; and Zhao, D. 2019b. Jointly Learning Entity and Relation Representations for Entity Alignment. In *EMNLP-IJCNLP*, 240–249.
- Xie, R.; Liu, Z.; Luan, H.; and Sun, M. 2017. Image-embodied knowledge representation learning. In *IJCAI*, 3140–3146.
- Xiong, W.; Yu, M.; Chang, S.; Guo, X.; and Wang, W. Y. 2018. One-Shot Relational Learning for Knowledge Graphs. In *EMNLP*, 1980–1990.
- Xu, K.; Wang, L.; Yu, M.; Feng, Y.; Song, Y.; Wang, Z.; and Yu, D. 2019. Cross-lingual Knowledge Graph Alignment via Graph Matching Neural Network. In *ACL*, 3156–3161.
- Yang, H.-W.; Zou, Y.; Shi, P.; Lu, W.; Lin, J.; and Xu, S. 2019. Aligning Cross-Lingual Entities with Multi-Aspect Information. In *EMNLP-IJCNLP*, 4422–4432.
- Yang, K.; Liu, S.; Zhao, J.; Wang, Y.; and Xie, B. 2020. COTSAE: CO-Training of Structure and Attribute Embeddings for Entity Alignment. In *AAAI*. AAAI Press.
- Zhang, Q.; Sun, Z.; Hu, W.; Chen, M.; Guo, L.; and Qu, Y. 2019. Multi-view knowledge graph embedding for entity alignment. In *IJCAI*, 5429–5435. AAAI Press.
- Zhou, B.; Lapedriza, A.; Khosla, A.; Oliva, A.; and Torralba, A. 2017. Places: A 10 million Image Database for Scene Recognition. *TPAMI*.
- Zhu, Q.; Zhou, X.; Wu, J.; Tan, J.; and Guo, L. 2019. Neighborhood-aware attentional representation for multilingual knowledge graphs. In *IJCAI*, 10–16.
- Zhuang, Y.; Li, G.; Zhong, Z.; and Feng, J. 2017. Hike: A hybrid human-machine method for entity alignment in large-scale knowledge bases. In *CIKM*, 1917–1926.